

(*) IO virtualization for VMs

~ system-view of ~~modern~~ OS includes devices.

~ device (external devices)



(*) device drivers use the interface to communicate w/ devices.
 - OS uses the device driver to translate actions from its own subsystems.

at least 3 ways for the OS to communicate ...

(i) PIO — programmed IO.

↳ special ISA instructions to deal w/ IO.

eg. in, out _____ serial port.

↳ port index, data

(ii) MMIO ~ memory mapped IO

↳ phy-mem is mapped to device interface state.

↳ r/w mem \Rightarrow r/w to device regs.

(iii) DMA ~ where devices. r/w directly to/from memory.

↳ setup by instruction on CPU by the OS.

Quiz #1

3rd Feb
in-class

[memory]

[live migration] [post copy]

[ramvs] [insight]

[snowflake]

e.g.: consider a disk, with following interface

reg x : sector number

y : R/W

z : VA (memory address)

in/out : instr. to state purpose

$x \leftarrow 10$

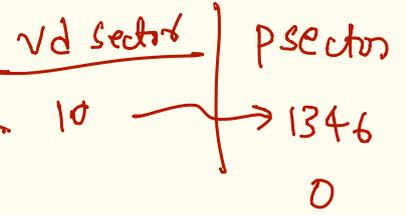
$y \leftarrow 0$

$z \leftarrow kva$

in

read from sector 10

& copy to kva specified in reg. z.

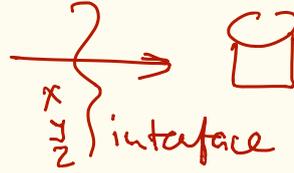


* Kva hyp.

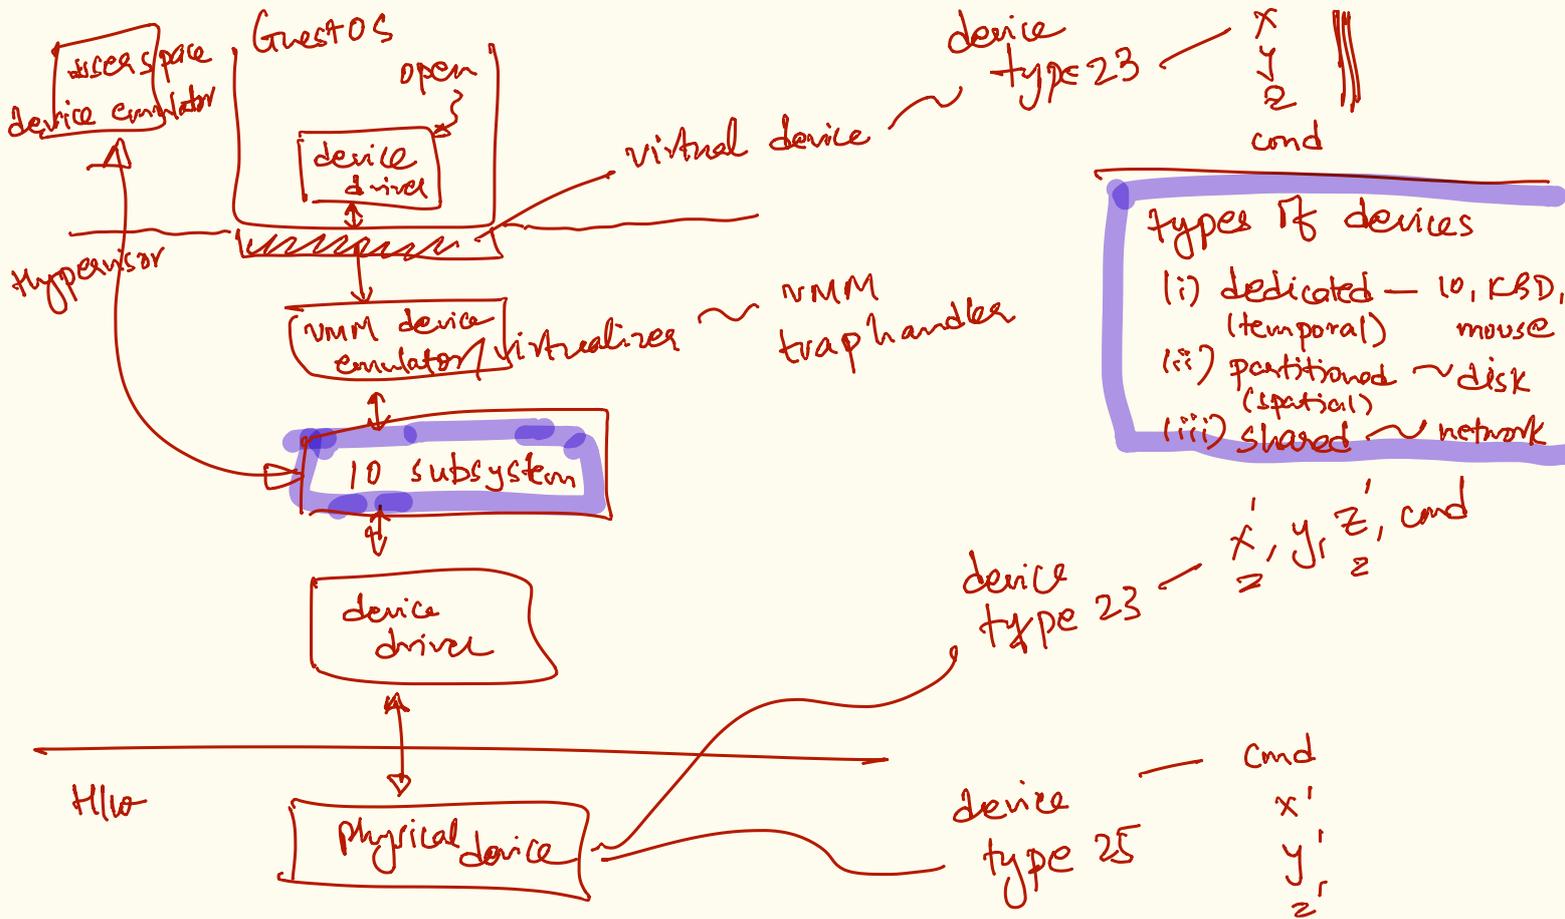
privileged instr.

fdopen → open

device block num → device drivers



Generalized IO virtualization view with VMs.



- types of devices
- (i) dedicated — IO, KBD, mouse (temporal)
 - (ii) partitioned ~ disk (spatial)
 - (iii) shared ~ network

* virtual device may/may not be the same as the physical device. e.g. vdisk ← file

* virtual device maybe virtualized via diff. abstractions.

① Full device virtualization (fully virtualized VMs).

~ guest OS does not know about the VMM
 ~ VMM virtualizes and handles/emulates the complete device interface.

device type 23

X
 Y
 Z
 Cmd. ← (i) privileged instr. generates trap.

~
 ⊗ actual eg - e1000 NIC
 ⊗ 2 complete IO stacks.

(ii) patch instr. to gen. trap.
 (iii) write-protect mmio pages to generate trap.

② PV IO virtualization — split device (driver) model.

device type 23

X ← IO trap →
 Y ← 0 trap →
 Z ← Kva trap →
 Cmdregs ← 0/1 trap →

PV

} → single hypercall
 hc-device IO (X, Y, Z, Cmd, devid)

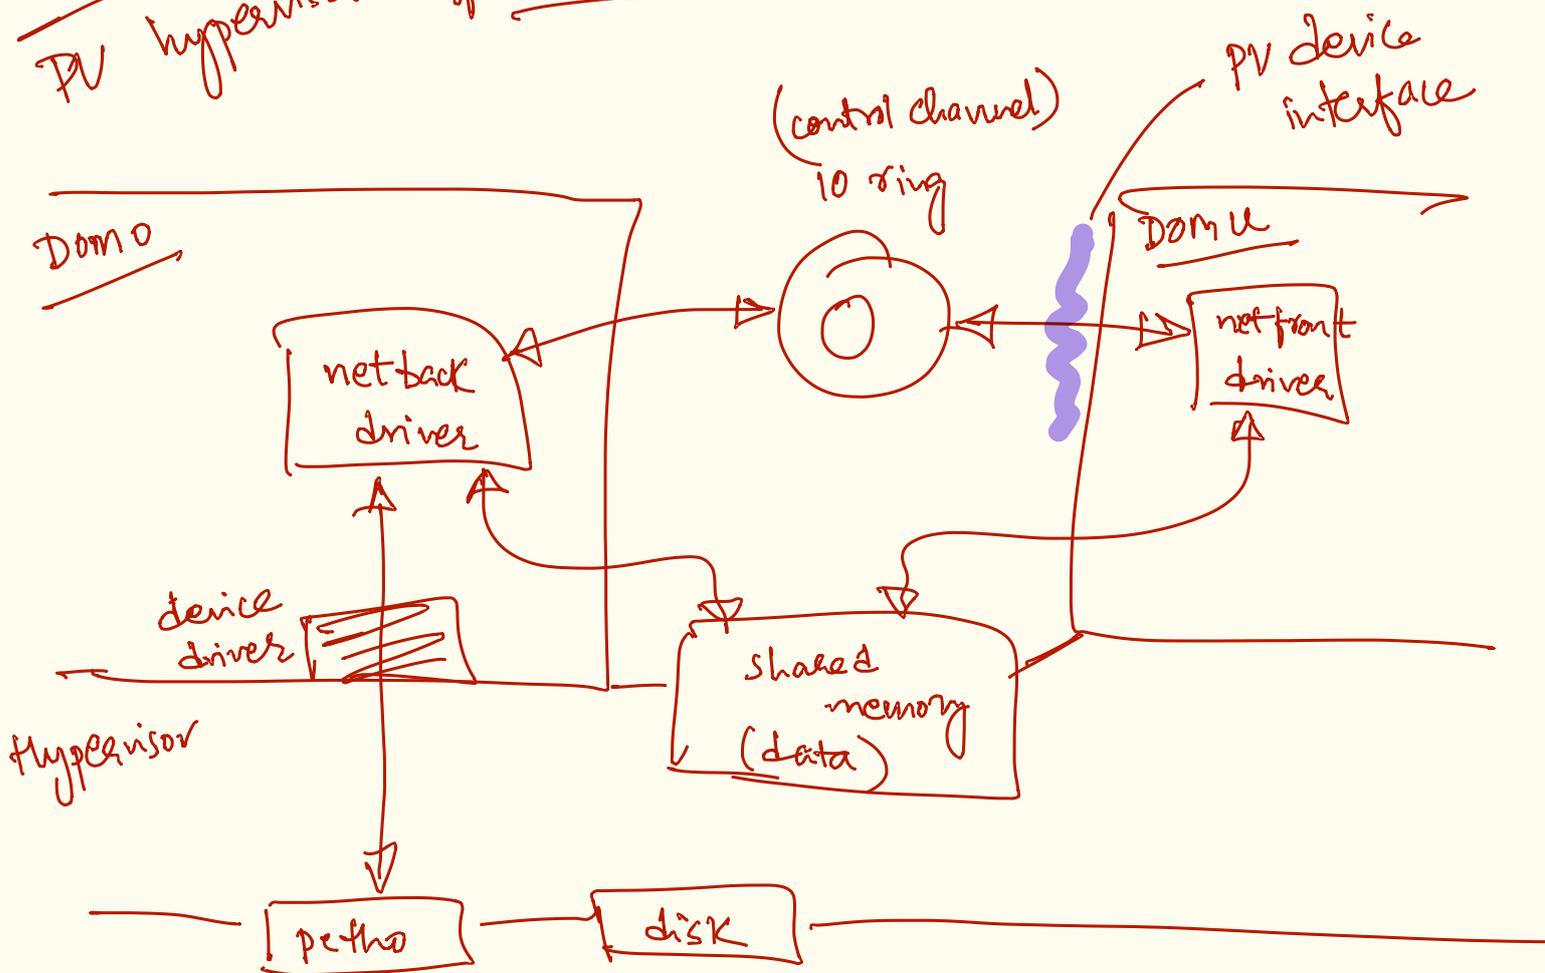
⊗ change device interface to a virtual interface
 non-physical/non-real

⇒ send all control+data info in a single hypercall.

⇒ pack multiple units of IO in a single call.

→ simplify device specification for greater efficiency!

Xen // virtio (KVM, VMware) ..
 PV hypervisor



Q1. How deep should the n/w stack be on the guest OS?

Q2. How does Dom0 get/receive pkts destined for DomU?